

## What is claimed is:

[Claim 1] 1. A multi-processor system comprising:

a plurality of snoop tag partitions for storing snoop entries;  
a plurality of external interconnect buses coupled between a plurality of snoop tag partitions, the plurality of external interconnect buses carrying cache coherency requests that include a snoop address;  
a plurality of memory controllers coupled to a shared main memory;  
a plurality of local processors for executing instructions and reading and writing data;  
a plurality of local caches, coupled to the plurality of local processors, for storing cache entries that contain instructions or data used by the plurality of processors;  
internal interconnect buses that couple the plurality of snoop tag partitions to the plurality of local caches and to the plurality of external interconnect buses;  
wherein each snoop tag partition in the plurality of snoop tag partitions contains snoop entries arranged into snoop sets, wherein a snoop index selects one of the snoop sets as a selected snoop set, wherein all snoop entries within a snoop set have a same snoop index but are able to have different snoop tags;  
wherein each local cache in the plurality of local caches contain cache entries arranged as multi-way cache sets, wherein a cache index selects one of the cache sets as a selected cache set, wherein all cache entries within a cache set have a same cache index but have different cache tags;  
wherein the snoop address carried over the internal interconnect buses comprises a tag portion for matching with a cache tag, a cache-index portion having the cache index for selecting the selected cache set, and an offset portion of data within a selected cache entry, wherein the cache-index portion further comprises a snoop-index portion having the snoop index for selecting the selected snoop set, a chip-select portion, and an interleave portion;  
wherein the chip-select portion of the snoop address selects a selected group of snoop tag partitions in the plurality of snoop tag partitions;

wherein the interleave portion of the snoop address selects a selected snoop tag partition in the plurality of snoop tag partitions within the selected group of snoop tag partitions;

wherein the selected snoop tag partition responds to the cache coherency request having the snoop address and stores a snoop tag in a snoop entry within the selected snoop set selected by the snoop index;

wherein other snoop tag partitions do not respond to the cache coherency request,

wherein the selected snoop tag partition is selected by the chip-select portion and the interleave portion of the snoop address which are subsets of the cache index,

whereby processing of snoop requests are partitioned across the plurality of snoop tag partitions by the chip-select portion of the snoop address.

[Claim 2] 2. A multi-processor system comprising:

a plurality of clusters coupled to a shared main memory;

a plurality of external interconnect buses coupled between the plurality of clusters, the plurality of external interconnect buses carrying cache coherency requests that include a snoop address;

wherein each of the plurality of clusters comprises:

a plurality of memory controllers coupled to the shared main memory;

a plurality of local processors for executing instructions and reading and writing data;

a plurality of local caches, coupled to the plurality of local processors, for storing cache entries that contain instructions or data used by the plurality of processors;

a plurality of snoop tag partitions for storing snoop entries;

internal interconnect buses that couple the plurality of snoop tag partitions to the plurality of local caches and to the plurality of external interconnect buses;

wherein each snoop tag partition in the plurality of snoop tag partitions contains snoop entries arranged into snoop sets, wherein a snoop index selects one of the snoop sets as a selected snoop set, wherein all snoop

entries within a snoop set have a same snoop index but are able to have different snoop tags;  
wherein each local cache in the plurality of local caches contain cache entries arranged as multi-way cache sets, wherein a cache index selects one of the cache sets as a selected cache set, wherein all cache entries within a cache set have a same cache index but have different cache tags;  
wherein the snoop address carried over the internal interconnect buses comprises a tag portion for matching with a cache tag, a cache-index portion having the cache index for selecting the selected cache set, and an offset portion of data within a selected cache entry, wherein the cache-index portion further comprises a snoop-index portion having the snoop index for selecting the selected snoop set, a chip-select portion, and an interleave portion;  
wherein the chip-select portion of the snoop address selects a selected cluster in the plurality of clusters;  
wherein the interleave portion of the snoop address selects a selected snoop tag partition in the plurality of snoop tag partitions within the selected cluster;  
wherein the selected snoop tag partition responds to the cache coherency request having the snoop address and stores a snoop tag in a snoop entry within the selected snoop set selected by the snoop index;  
wherein other snoop tag partitions do not respond to the cache coherency request,  
wherein the selected snoop tag partition is selected by the chip-select portion and the interleave portion of the snoop address which are subsets of the cache index,  
whereby processing of snoop requests are partitioned across the plurality of clusters by the chip-select portion of the snoop address.

[Claim 3] 3. The multi-processor system of claim 2 wherein each snoop entry stores a snoop tag;

wherein each cache entry stores a cache tag;

wherein the tag portion of the snoop address has a same number of address bits as the snoop tag and as the cache tag,

wherein the tag portion matches the cache tag of the selected cache entry;  
wherein the tag portion is matched with the snoop tag of a selected snoop entry in the selected snoop set,  
whereby cache tags and snoop tags are a same size and interchangeable.

[Claim 4] 4. The multi-processor system of claim 2 wherein a number of the snoop entries per snoop set in one snoop tag partition in the plurality of snoop tag partitions is equal to a number of the cache entries per cache set multiplied by a number of local caches in the plurality of local caches multiplied by a number of clusters in the plurality of clusters,  
wherein a total number of the snoop entries in the multi-processor system equals a total number of cache entries in all of the local caches in the plurality of local caches in the multi-processor system.

[Claim 5] 5. The multi-processor system of claim 2 wherein the plurality of clusters comprises N cluster chips, wherein N is a whole number;  
wherein the multi-processor system is expandable by adding additional cluster chips to the multi-processor system and by increasing a number of address bits in the chip-select portion and decreasing a number of address bits in the snoop-index portion of the snoop address.

[Claim 6] 6. The multi-processor system of claim 5 wherein N is expandable from 1 to 16 and wherein the plurality of local caches comprises M local caches for each cluster;  
wherein each cache set comprises W cache entries;  
wherein each snoop set comprises Q snoop entries, wherein Q is equal to  $N * M * W$ ;  
wherein M, W, and Q are whole numbers.

[Claim 7] 7. The multi-processor system of claim 6 wherein W is 4 or 8,  
wherein each local cache is a 4-way or an 8-way set-associative cache.

[Claim 8] 8. The multi-processor system of claim 7 wherein the plurality of snoop tag partitions comprises  $S$  snoop tag partitions on each cluster; wherein a number of address bits in the interleave portion of the snoop address is  $B$ , wherein  $S$  is equal to  $2^B$ , wherein  $S$  and  $B$  are whole numbers.

[Claim 9] 9. The multi-processor system of claim 8 wherein a number of address bits in the chip-select portion of the snoop address is  $C$ , wherein  $N$  is equal to  $2^C$ , wherein  $C$  and  $N$  are whole numbers.

[Claim 10] 10. The multi-processor system of claim 9 wherein a number of address bits in the cache-index portion of the snoop address is  $D$ , wherein there are  $2^D$  cache sets in each local cache; wherein a number of address bits in the snoop-index portion of the snoop address is  $F$ , wherein there are  $2^F$  snoop sets in each snoop tag partition; wherein  $D$  is equal to  $F+B+C$ , wherein  $B$ ,  $C$ ,  $D$ , and  $F$  are whole numbers.

[Claim 11] 11. The multi-processor system of claim 10 wherein  $S$  is 2,  $M$  is 3,  $W$  is 4, and  $N$  is 16; wherein each snoop set in each snoop tag partition has 192 snoop entries.

[Claim 12] 12. The multi-processor system of claim 8 wherein each memory controller in the plurality of memory controllers is tightly coupled to a snoop tag partition in the plurality of snoop tag partitions, wherein snoop addresses mapping to a snoop tag partition by the chip-select and interleave portions of the snoop address have memory accesses processed by a memory controller tightly coupled to the snoop tag partition,

wherein each memory controller is for accessing a partition of a memory space in the shared main memory, wherein the partition is  $1/(N*S)$  of the memory space.

[Claim 13] 13. The multi-processor system of claim 12 wherein each snoop entry stores the snoop tag and a snoop state of a corresponding cache line having a cache tag matching the snoop tag;

wherein each cache entry stores the cache tag and a cache state of the corresponding cache line;

wherein the snoop state is invalid, shared, owner, or modified;

wherein the cache state is invalid, shared, owner, exclusive, or modified;

wherein the snoop state is encoded by two state bits, but the cache state is encoded by 3 state bits.

[Claim 14] 14. The multi-processor system of claim 13 wherein the snoop state is modified when the cache state is exclusive or the cache state is modified.

[Claim 15] 15. A coherent multi-chip multi-processor system comprising:

a main memory;

an external interconnect among cluster chips including a first cluster chip, a second cluster chip, a third cluster chip, and an Nth cluster chip;

wherein the first, second, third, and Nth cluster chip each comprise:

a first memory controller for accessing the main memory;

a second memory controller for accessing the main memory;

a snoop interconnect coupled to the external interconnect;

a first snoop tag partition, coupled to the first memory controller and to the snoop interconnect, for storing snoop entries in snoop sets selected by a snoop index in an address, each snoop entry for storing a snoop tag from the address or that matches a tag portion of the address;

a second snoop tag partition, coupled to the second memory controller and to the snoop interconnect, for storing snoop entries in snoop sets selected by the snoop index in an address, each snoop entry for storing a snoop tag from the address or that matches the tag portion of the address;

wherein each snoop set identified by a corresponding snoop index stores snoop entries for all caches in the coherent multi-chip multi-processor system for a corresponding cache index that has the corresponding snoop index as a subset;

wherein a cache index in the address contains an interleave bit that is not in the snoop index, the interleave bit selecting either the first snoop tag partition or the second snoop tag partition for storing the snoop entry for the address;

a first processor for executing instructions;

a second processor for executing instructions;

a first local cache, coupled between the first processor and the snoop interconnect, for storing cache entries that each store a cache tag and data; and

a second local cache, coupled between the second processor and the snoop interconnect, for storing cache entries that each store a cache tag and data; wherein the cache entries are arranged into cache sets each having at least four associative cache entries with a same cache index but different cache tags and data, wherein the cache index selects a cache set in a local cache;

wherein the address comprises a tag portion for matching with or storing as the snoop tag and as the cache tag, and the cache index;

wherein the cache index in the address comprises the snoop index, chip-select bits, and the interleave bit;

wherein the chip-select bits in the address select the first and second snoop tag partitions in either the first cluster chip, the second cluster chip, the third cluster chip, or the Nth cluster chip for storing the snoop entry for the address,

whereby the chip-select bits in the address select a cluster chip while the interleave bit selects a snoop tag partition for storing the snoop entry for the address.

[Claim 16] 16. The coherent multi-chip multi-processor system of claim 15 wherein each cluster chip further comprises:

a third processor for executing instructions;  
a third local cache, coupled between the third processor and the snoop interconnect, for storing cache entries that each store a cache tag and data;  
wherein each snoop set identified by a corresponding snoop index stores at least  $N \times 12$  snoop entries, wherein  $N$  is a whole number indicating a number of cluster chips in the coherent multi-chip multi-processor system.

[Claim 17] 17. The coherent multi-chip multi-processor system of claim 15 wherein a number of chip-select bits is  $B$ , wherein  $B$  is a whole number, and wherein  $N$  is  $2^B$  cluster chips;

wherein the cache index has  $B+1$  more address bits than the snoop index.

[Claim 18] 18. A processing cluster chip for a multiprocessing system comprising:

first memory controller means for accessing a memory;  
first snoop tag partition means, coupled to the first memory controller means, for storing snoop tags in snoop entries arranged in snoop sets selected by a snoop index of an address;  
second memory controller means for accessing a memory;  
second snoop tag partition means, coupled to the second memory controller means, for storing snoop tags in snoop entries arranged in snoop sets selected by a snoop index of an address;  
interconnect means for connecting the first and second snoop tag partitions to caches on other processing cluster chips and to local caches;  
first processor means for executing programmable instructions;  
first cache means, between the interconnect means and the first processor means, for storing data from the memory in cache entries arranged in cache sets selected by a cache index of the address;  
second processor means for executing programmable instructions;



second cache means, between the interconnect means and the second processor means, for storing data from the memory in cache entries arranged in cache sets selected by a cache index of the address,  
third processor means for executing programmable instructions; and  
third cache means, between the interconnect means and the third processor means, for storing data from the memory in cache entries arranged in cache sets selected by a cache index of the address,  
wherein each cache set contains 4 cache entries and each cache entry contains a cache tag, a dirty bit, and data;  
wherein the address contains tag bits forming the cache tag or the snoop tag, and a cache index for selecting the cache set;  
wherein the cache index contains the snoop index and chip-select bits and an interleave bit;  
wherein the interleave bit in the address selects the first snoop tag partition means for processing a request using the address when the interleave bit is in a first state, and the interleave bit in the address selects the second snoop tag partition means for processing the request using the address when the interleave bit is in a second state;  
wherein the chip-select bits in the address select a processing cluster chip that contains a selected snoop tag partition means for processing the request using the address;  
whereby the chip-select and interleave bits in the address select the first or second snoop tag partition means and the processing cluster chip for processing the request.

**[Claim 19]** 19. The processing cluster chip for a multiprocessing system of claim 18 wherein the multiprocessing system comprises N processing cluster chips;

wherein each snoop set comprises  $N*4*3$  snoop entries, one snoop entry for each cache entry for a selected cache index for all processing cluster chips in the multiprocessing system.

[Claim 20] 20. The processing cluster chip of claim 19 wherein each cache entry stores a 3-bit state of the data, while each snoop entry stores a 2-bit state of the data stored in the cache entry.